

Amendments to the Specification:

On page 18, please replace paragraph [0035] with the following paragraph where additions are underlined and deletions are in strike-out font:

-- After the entire file has been read and stored in physical memory, the process 400 can fulfill subsequent file system read requests relating to that file from memory without having to access the device. Consequently, the process 400 optionally can deactivate the device immediately after reading the entire file from the device in block 422. In any event, the next time that a file system read request is intercepted seeking a portion of the cached file (402) and the process 400 confirms that the entire file has been loaded into physical memory (404), the process 400 can return the requested file segment from memory (406) without having to activate or access the device.--.

Beginning on page 19, please replace paragraph [0037] with the following paragraph where additions are underlined and deletions are in strike-out font:

-- For example, assume that the intercepted read request (402) is a request to read a row from a table in a 300 MegaByte (MB) relational database file. In that case, the entire 300 MB file generally is too large to be read entirely and stored in a typical computer system's physical memory. Instead, the process 400 may read in not only the requested row from the database file, but also an additional portion of the database file that is likely to be accessed by future read requests, in block 418. The amount of the file to be read into physical memory may depend on one or more factors, including available physical memory space, user defined parameters, file size and the like, as may be determined by various rules, in block 418A. The additional portions of the file to be read into physical memory may be chosen intelligently, for example, using predictive caching techniques or the like. In the database file example, the additional portions to be read (i.e., the superset of the requested file portion) may include adjacent rows, the entire table, and/or other logically related information fields. In this regard, a file access monitor

process could be implemented that keeps track of, and identifies trends relating to, which files, and which portions of files, have been accessed recently, e.g., for the last 3-4 days or so. The intelligent decisions as to which file portions to be read into physical memory by process 400 could be based on the information collected and maintained by the file access monitor, in block 418A. Alternatively, or in addition, a secondary memory paging manager (discussed below) could use information from the file access monitor to make intelligent decisions about which page or pages are to be removed from physical memory, e.g., in order to accommodate additional portions of a file read in response to an intercepted read request.--.